Adaptive Child-Robot Interaction: Enhancing Educational Outcomes through Applied Behavior Analysis and Real-Time Replanning

Giulia Berettieri[1,†], Anna Allegra Bixio[1,*,†], Lucrezia Grassi[1], Carmine Recchiuto[1] and Antonio Sgorbissa[1]

[1]DIBRIS Department, University of Genoa, Genoa, Italy

Abstract

This paper introduces a robotic system architecture for behavior management in educational settings, inspired by Applied Behavior Analysis (ABA). The system utilizes a socially assistive robot (SAR) that autonomously delivers educational activities, monitors behavior, and proposes management strategies using the Planning Domain Definition Language (PDDL) framework. By integrating cloud-based processing with local execution, the robot adapts in real-time to enhance social interactions. The study outlines future works, including an experimental protocol designed to enhance the management of children's behavior during child-robot interactions.

Keywords

Behavior Management, Real-Time Adaptation, Child-Robot Interaction

1. Introduction

This paper proposes a novel robotic architecture for behavior management in educational settings, inspired by Applied Behavior Analysis (ABA), an evidence-based approach that aims at explaining how behavior works and is affected by the environment [1]. Despite its controversial role in supporting individuals with Autism Spectrum Disorder (ASD) [2][3], we think that an ABA-inspired robot capable of reasoning about the purpose of others' behaviors can be a useful support in educational settings. ABA is presently used in schools due to its effective behavior management techniques [4] [5]. Applying these strategies to social robotics could improve the quality of child-robot interaction (CRI) in learning environments.

Socially assistive robots (SARs) are increasingly employed in education to create personalized interventions [6] and computational techniques are used to help tailor learning activities to the capabilities of the student [7][8]. Current research on robot-mediated interventions using ABA principles primarily focuses on supporting children with ASD. SARs are expected to act as social mediators, employing typical ABA methods to enhance individuals' performance in social skills [9][10][11].

Recently, research on behavior change strategies in educational robotics settings has gained significant attention. Studies have demonstrated that a motivational and persuasive approach of the robot, which involves the provision of tailored feedback and individualized goals, has a positive impact on encouraging children to follow a healthier lifestyle [12][13]. Furthermore, [14] showed that the design of the robot itself can significantly influence its effectiveness.

ALTRUIST, BAILAR, SCRITA, WARN 2024: Workshop on sociAL roboTs for peRsonalized, continUous and adaptIve aSsisTance, Workshop on Behavior Adaptation and Learning for Assistive Robotics, Workshop on Trust, Acceptance and Social Cues in Human-Robot Interaction, and Workshop on Weighing the benefits of Autonomous Robot persoNalisation. August 26, 2024, Pasadena, USA

*Corresponding author.

†These authors contributed equally.

s4816056@studenti.unige.it (G. Berettieri); annaallegra.bixio@edu.unige.it (A. A. Bixio); lucrezia.grassi@edu.unige.it (L. Grassi); carmine.recchiuto@dibris.unige.it (C. Recchiuto); antonio.sgorbissa@unige.it (A. Sgorbissa)

https://orcid.org/0009-0002-4569-8871 (A. A. Bixio); https://orcid.org/0000-0001-6363-3962 (L. Grassi); https://orcid.org/0000-0001-9550-3740 (C. Recchiuto); https://orcid.org/0000-0001-7789-4311 (A. Sgorbissa)

Additionally, studies have found evidence that managing behavior in HRI interventions requires robotic systems with user cognitive profiling [15]. Research in this domain examines how Theory of Mind (ToM) – the ability to recognize others' mental states [16] – can be replicated computationally. For instance, some studies propose models that enable robots to infer goals [17], while others develop neural networks to predict actions based on behavior [18]. Moreover, cognitive architectures inspired by existing theories provide social robots with artificial emotions to interpret social environments [19]. Several studies investigate how ToM skills, such as the ability to understand false-belief tasks, influence human trust and decision-making [20][21]. Consequently, the interest in implementing a robotic architecture that assumes others' mental states have inspired research on the role of deception in HRI [22]. These findings highlight that robotic ToM abilities affect human tendencies to attribute mental states to social robots, enhancing their acceptance as trustworthy companions.

The contributions of this work are the following:

- We developed a system architecture utilizing the PDDL framework, enabling real-time adaptation and allowing the robot to autonomously manage interactions.
- We established an experimental protocol designed to enhance the management of children's behavior during child-robot interactions.

The article is structured as follows. Section 2 introduces the ABA methods and describes the scenario. Section 3 details the system's architecture and behavioral planning. Section 4 presents a case study on behavior function recognition. Section 5 summarizes ongoing performance evaluation experiments. Section 6 outlines the proposed experimental protocol. Finally, Section 7 discusses the conclusions and limitations.

2. Background and problem statement

Applied Behavior Analysis (ABA) is an evidence-based approach aimed at improving socially significant behaviors [1]. ABA identifies the environmental variables responsible for behavior change through a validated approach that involves the analysis of factors influencing behavior and the application of scientific inquiry methods. Essentially, ABA involves a behavioral assessment to define and quantify target behaviors, as well as to identify desired achievements to ensure the generalization and maintenance of behavior change [1]. In the ABA framework, behavior is modeled as purposeful (functional) and influenced by stimuli occurring before (antecedents) and after the behavior (consequences). Some consequences could increase the likelihood of a behavior to occur and are known as reinforcements. The process of gathering information about the relationship between behavior, its antecedent and consequence, to understand the behavior's purpose, is called Functional Behavior Assessment (FBA) [23]. The functions of behavior, described as purposes of behaviors, can be divided into four types [1]:

- *Gain Attention*, i.e. the purpose of behaviors aiming at receiving attention from others.
- *Gain a Tangible*, i.e. the purpose of behaviors aiming at having access to materials or preferred stimuli.
- *Gain or Avoid Sensory Stimulation*: i.e. the purpose of behaviors aiming at producing or avoiding a sensory stimulation.
- *Escape*, i.e. the purpose of behaviors aiming at avoiding unpleasant situations, such as attention, demanding tasks, or boring activities.

To detect the function, an assessment of the antecedent and consequence of the target behavior is conducted. Teachers frequently face behaviors that interfere significantly with students' and their peers' learning. ABA assessment methods are currently implemented in schools [4][24][5], and FBA can support teachers in providing a clear method to manage behaviors complex to handle. Conversely, robots as tutors may offer individual attention by adapting their behavior to support a single child's

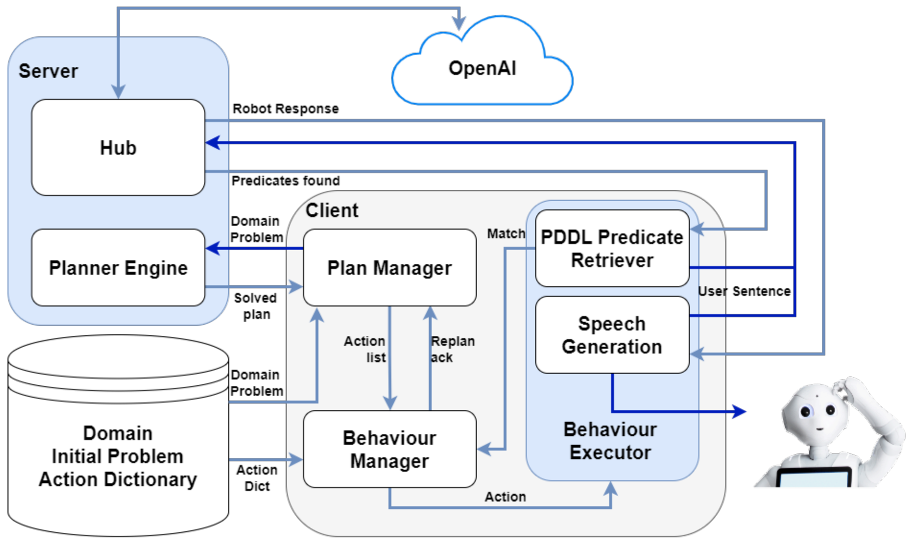

Figure 1: System architecture integrating cloud services with local execution.

learning, without disrupting classroom activities. In this context, SARs could be valuable tools to alleviate the excessive time spent on classroom management.

Given these premises, the long-term objective of this work is to develop an autonomous social robot capable of planning educational activities, continuously monitoring children's actions, conducting FBA to identify the purposes of challenging behaviors, and proposing alternative strategies to handle them. Our aim is not to categorize behaviors but to address a common classroom struggle: managing behaviors considered "challenging" because they present major challenges to maintaining a productive learning environment. A robot that can tailor learning interactions and respond effectively to others' behaviors through the use of effective strategies could greatly improve the integration of robotics in education.

3. System architecture and planning for behavioral change

To achieve effective and autonomous behavior management in CRI, we have designed a software architecture that integrates cloud-based processing with local execution capabilities.

Figure 1 illustrates the detailed layout of the proposed system architecture, which is divided into two main components: the Server and the Client. Each component plays a crucial role in the system's overall functionality. The server hosts two primary elements: the Hub and the Planner Engine. The Hub manages connections to the external cloud service OpenAI, facilitating the integration of large language models (LLMs). The Planner Engine leverages Fast Downward, a highly efficient planning system based on the Problem Domain Definition Language (PDDL). PDDL operates with a domain file, that defines the available predicates to describe the world and actions that can modify its state, and a problem file, that specifies the initial state and the desired goal state to be achieved. The Planner Engine processes inputs from the client, specifically the domain PDDL file and the initial problem file. These files outline the available actions, their effects, preconditions, and the initial state of the problem that the robot needs to address. For example, in our scenario, the plan might involve playing a memory card game with a child, including actions such as explaining the rules and providing feedback.

The client, embedded within the robot, comes preloaded with the necessary domain and initial problem files. The Plan Manager is responsible for sending these files to the Planner Engine and retrieving the generated plan. It acts as a bridge between the Planner Engine and the rest of the system, ensuring that the plan is accurately interpreted and executed, even if replanning is required.

Upon receiving the plan, the Behaviour Manager takes charge, iterating through the plan and determining the actions that need to be executed by checking the predicates.

As actions are performed, the PDDL Predicate Retriever asserts new predicates about the current state of the world based on the interaction with the child and the environment. The goal is to monitor the plan advancement, to ensure that the asserted preconditions and effects associated with each performed action, periodically retrieved through observations of the world, align with the expected preconditions and effects as they have been previously planned (this will be handled by the Behavior Manager, as explained in the following). To achieve this, the robot uses a multi-sensor approach to recognize behaviors.

On one hand, the robot employs its sensors to make assertions about the current state of the world; on the other hand, it attempts to infer the child's mental state through conversation. To this end, the robot employs Microsoft Azure services to transcribe the child's spoken words captured by microphones. The transcribed text is then sent to OpenAI, which analyzes the content and context to generate appropriate responses and identify predicates related to the child's behavior.

It is important to explain how the mechanism of predicate retrieval works. When the server loads the first domain, it is given a list of relevant predicates referring to the agent. Thus, when an action begins, all predicates unrelated to the agent's state are accounted for. OpenAI retrieves a predicate if it is not already on the list, or, if it has been mentioned, it will not be removed unless explicitly negated in the dialogue. For example, if a user mentions that they are performing a task, the predicate "doing" is registered. Even if the user talks about something else, this does not imply that they have stopped doing the task, so the predicate is neither removed nor negated. This architecture ensures an adaptable and efficient task execution system, allowing the robot to offload complex planning and language processing tasks to the server while focusing on real-time interaction and action execution.

The Behavior Manager compares asserted predicates with the expected preconditions and effects in the plan. If the predicates do not match, the Behavior Manager triggers the Plan Manager, which may then initiate a replan. This process follows a "match and go" approach, where the robot continuously evaluates the environment against a set of predefined predicates and conditions that must be met for an action to be executed. When a new situation arises, the robot checks its current state against the effects of the current action and all other actions' preconditions. If these conditions "match" the expected criteria, the robot can either proceed to execute the next action — hence the term "go" — or request a replan. This method ensures that the robot's actions are contextually adaptive.

4. A case study: *Escape* and *Gain a Tangible*

Drawing from ABA and FBA principles, two behavior functions were identified to be recognized for preliminary testing of the model's validity: *Gain a Tangible* and *Escape*. The aim is that, through verbal and non-verbal interaction with the child, the robot would be able to autonomously recognize the occurrence of the following behaviors:

- *Gain a Tangible*: When pleasant activities are interrupted by time limits, negative reactions of the child may be elicited to prolong the activity.
- *Escape*: When presented with a task, a child may perceive it as too difficult or too boring, potentially provoking adverse responses to avoid the task.

With a comprehensive understanding of these psychological approaches, we developed a PDDL domain capable of autonomously managing various relevant scenarios where the two behavior functions mentioned above can influence a child's behavior. Table 1 details the actions that have been modeled. Notably, the PDDL domain includes additional actions beyond behavior management, such as moving the robot to different areas or greeting a parent arriving early to pick up their child.

The list of available PDDL actions includes those that enable the robot to apply the appropriate intervention strategy once a behavioral function has been detected, such as `StrategyWant2Play` for addressing *Gain a Tangible*, and `StrategyBoringTask` or `StrategyHardTask` for addressing *Escape*. However, before implementing a strategy, it is necessary to better understand the child's mental state

Action	Description
`DiscoverToM`	Converse with child to guess their mental state
`GoodJob`	Praise child for achievement
`GoodJobStrategy`	Praise child for achievement with strategy
`PresentTask`	Explain task rules in detail
`PutAway`	Ensure the child stops playing
`StrategyWant2Play`	Apply strategy if child still wants to play
`StrategyBoringTask`	Apply strategy if child finds task boring
`StrategyHardTask`	Apply strategy if child finds task hard
`Wait`	Wait and monitor task execution
`WarningCleanUp`	Warn about time and monitor completion
`ConfirmWant2Play1`	Confirm *Gain a Tangible* function
`ConfirmWant2Play2`	Confirm *Gain a Tangible* function
`ConfirmHardTask`	Confirm *Escape* function
`ConfirmBoringTask`	Confirm *Escape* function

Table 1
PDDL Actions related to Behavior Management.

and confirm that the initial assumptions about the detected functions are correct. This is achieved using the action `DiscoverToM` to initiate a conversation with the child and infer their mental state.

For a thorough understanding of the process, consider the following plan that is produced in a card game scenario.

```
(Interact game kid memory)
(PresentTask game kid memory)
(Wait game kid memory)
(WarningCleanUp game kid memory)
(PutAway game kid memory)
(ConfirmAllGood game kid memory x1 x2)
(GoodJob game kid memory x1 x2)
```

The initial plan involves a brief interaction with the child (`Interact ...`), followed by the robot explaining the rules of the memory card game (`PresentTask ...`). As the child begins playing, the robot waits and monitors their engagement, occasionally inquiring about their current state (`Wait ...`). Shortly before the designated playtime ends, the robot notifies the child that their time is almost up (`WarningCleanUp ...`). When the time is up, the robot instructs the child to stop playing and put away the cards (`PutAway ...`), ensures the child is no longer playing (`ConfirmAllGood`), and finally, praises the child and says goodbye (`GoodJob`). Actions are detailed in Table 1.

During the second action (`Wait ...`), the robot occasionally interacts with children through OpenAI-supported verbal exchanges. As already mentioned, this is achieved using a carefully designed prompt that explicitly queries OpenAI to ground predicates based on the children's responses. The list of predicates to be grounded through conversation includes those that describe the child's activity as well as the child's mental state regarding the task (the ability of OpenAI to demonstrate a Theory of Mind (ToM) has been previously discussed [25]), which serves as the basis for identifying task-related issues by directly querying the children.

The paragraphs below present some examples of verbal interaction logs related to this situation. User inputs are denoted by "U," the robot's responses by "R," and the retrieved predicates by "P."

...
U: Pepper, can we talk while I am playing?
R: Giulia, let's play together!
P: `warning, not_waiting, want_to_play`
U: No, don't worry, I don't want to play anymore.
R: Giulia, why are you so upset?
P: `warning, not_waiting, want_to_play,`

`not_agreed`, `not_doing`
Replan is needed

In the exchange above, the robot detects a shift in the child's engagement, recognizing that Giulia is no longer interested in continuing the game. This prompts replanning starting from the current state (see Section 3), with the asserted predicate `not_doing` acting as a trigger. Since this predicate does not match the preconditions of the `WarningCleanUp` action, which was originally expected as the next step, but matches the preconditions of the alternative action `DiscoverToM`, it signals the need for replanning.

In response to this interaction, a new plan is formulated, which includes the following action:

```
(DiscoverToM game kid memory).
```

Once again, this action will involve verbal interaction with the child using OpenAI functionalities, aiming to ground predicates from a given list of possible options based on the child's responses. The following log shows a possible outcome of the interaction.

...
U: Because this game is too difficult!
R: Giulia, what makes this game difficult for you?
P: `investigating, task_is_hard`
Replan is needed

After investigating the child's activity during the `(Wait ...)` action, the robot now explores the child's mental state in the `(DiscoverToM ...)` action. Based on the child's responses, the robot asserts the predicate `task_is_hard`. This update, performed by the PDDL Predicate Retriever, leads to the conclusion that the task is perceived as too challenging for the child, once again prompting the Plan Manager to initiate a replanning process. Specifically, the asserted predicates match the preconditions of the `ConfirmHardTask` action, indicating the behavioral function *Escape*, which necessitates the specific strategy `StrategyHardTask`. The newly retrieved plan is as follows:

```
(ConfirmHardTask game kid memory x1 x2)
(StrategyHardTask game kid memory x1 x2)
(Wait game kid memory)
(WarningCleanUp game kid memory)
(PutAway game kid memory)
(ConfirmAllGood game kid memory x1 x2)
(GoodJob game kid memory x1 x2)
```

This is just one possible scenario. After the strategy has been applied, other behavior functions may be detected. For instance, if the child disagrees with ending the task when the time is up, a new plan will be generated. In this second case, after confirming the behavior function *Gain a Tangible*, the resulting plan would be as follows:

```
(ConfirmWant2Play1 game kid memory x1 x2)
(StrategyWant2Play game kid memory x1 x2)
(GoodJobStrategy game kid memory x1 x2)
```

In this plan, it is first confirmed that the child still wants to play with the memory game, and the *Gain a Tangible* strategy is therefore applied. Precisely, the robot will grant an additional two minutes of play with a timer, and then it will ask the child to help clean up the cards.

5. Performance Evaluation

Following the framework outlined in Sections 3 and 4, we will conduct structured lab experiments focused on a memory game task, where participants match pairs of identical cards. These experiments aim to evaluate metrics relative to the robot's replanning abilities in controlled settings, using adult participants who simulate children's behavior.

Participants are briefed on the experimental protocol and engage in semi-scripted interactions with the robot. During these interactions, participants intentionally disrupt the robot's planned actions, prompting it to replan.

Key metrics being assessed include:

- *Replan Time and Frequency*: Measures the time from the introduction of a disruption to the completion of the replan and the frequency of replanning events, assessing the responsiveness of the replanning mechanism.
- *Cloud Response Time*: Evaluates the time taken for the robot to send a request, process cloud information, and receive a response, gauging the latency introduced by cloud dependency in the replanning process.
- *Efficacy of Replan*: Assesses the robot's ability to adjust actions, recognize predicates, and achieve intended outcomes despite interruptions.

These experiments provide insights into the robot's performance in adapting to real-time disruptions during interactions, paving the way for future experiments with actual children.

6. Experimental protocol

To assess our model's effectiveness in educational settings, we propose an experimental design comprising 20 child-robot interaction sessions. Education experts will retrospectively evaluate these sessions to assess the robot's accuracy in detecting behavior functions and the appropriateness of employed strategies. Our hypotheses posit that the functions identified by our model will align with experts' judgments, and the strategies employed will be approved by them.

6.1. Participants

The study involves a total of 30 participants, including 20 children aged 8-9 years randomly selected from Italian primary schools and 10 education experts, ideally primary school teachers. Each child engages in a one-on-one interaction with the robot and is accompanied by a caregiver who is asked to intervene during the study. Education experts assess child-robot interactions by reviewing their video recordings.

6.2. Hypotheses

Based on the evaluation of child-robot interactions conducted by education experts, the following hypotheses are proposed:

H1: There will be a significant agreement between the behavior functions assessed by experts and those detected by the robot during the experiment. A high level of agreement indicates the accuracy of the robot's capability to interpret the child's behavior.

H2: The strategies proposed by the robot during interactions will be positively evaluated by experts. A strong endorsement of the strategies employed would demonstrate increased validation of the robot's applicability in educational environments.

These hypotheses explore the alignment between education expert assessments and robot performance, contributing to understanding the applicability of robotics in education.

6.3. Experimental Procedure

During the experiment at the university laboratory, each child will independently interact with the humanoid robot Pepper, which operates based on the architecture detailed in Section 3. The robot autonomously plans the sequence of actions based on how the interaction evolves, following the outlined general rules.

Pepper starts by engaging the child and explaining the memory game, including the rules and timing. The child has 10 minutes to complete the game and is instructed to leave the room when the time is up. During the game, Pepper interacts minimally, asking only if the game is progressing well. If the robot detects that the task is too difficult, it allows the child to play with fewer cards. Conversely, if the game seems too boring, Pepper suggests playing with more cards to increase the challenge.

Two minutes before the time expires, Pepper provides a reminder. When the time is up, the caregiver is asked to call the child to leave. If the child complies, Pepper thanks them and says goodbye. If the child resists, Pepper recognizes the child's desire to continue playing and implements the following strategies: first, it grants an additional two minutes of play with a timer, and then it asks the child to help clean up the cards. Once the cards are put away, Pepper waves goodbye, and the participants leave the room. Each session, lasting approximately 15-20 minutes, is recorded.

We clarify that in the final phase of the experiment, caregiver intervention is intended to encourage the child to engage with the robot. We hypothesize that the presence of a familiar person may influence the child's noncompliance more than the robot's instructions. This allows for a more effective test of the system, as it increases the likelihood that the robot will need to implement strategies to manage the child's behavior. Moreover, Pepper interacts with minimal social cues to avoid influencing the child's task performance. Finally, the strategies implemented are carefully recommended by both the authors and the literature review. Specifically, one of the authors is an education specialist with expertise in classroom settings and reported the main strategies used in such contexts. Some of these, such as using a timer as a visual support, adjusting tasks based on students' entry skills, and providing praise after the expected behavior is performed, are well-documented in ABA and behavior management studies [4][26].

Approval from the Ethics Committee of the University of Genoa and parental consent will be obtained.

6.4. Measurements

For *H1*, we aim to evaluate the agreement between Pepper's identification of behavioral functions and experts' judgments across 20 recorded sessions. Two experts review each video to assess whether the child exhibits the functions *Gain a Tangible* and *Escape*. Videos include clickable commands for each function and an "End Function" command to signal the cessation of function detection. Experts provide judgments at any appropriate point during the videos. After collecting the responses, we calculate the degree of agreement between the two experts, followed by the degree of agreement between each expert and the robot. The mean of these results is then determined.

This data analysis is conducted for every 30-second segment of the video using a Weighted Cohen's K test. Each segment is labeled as either *No Function*, *Function 1*, or *Function 2* based on the buttons clicked by experts. The degree of agreement is computed and averaged across all segments, taking into account different combinations of responses between the two evaluators. For example, simultaneous identification of two distinct functions (*Function 1* and *Function 2*) is weighted more than when one rater identifies *Function 1* and the other identifies *No Function*.

For *H2*, we aim to assess the effectiveness of Pepper's strategies through expert evaluation. In the recorded videos where Pepper implements strategies, experts indicate which strategies they would choose in the same scenario. Afterward, experts complete a survey evaluating the acceptability, appropriateness, and feasibility [27] of Pepper's strategies.

This approach evaluates the accuracy of Pepper's behavioral function identification and expert perceptions of its strategies.

7. Conclusion

Our study provides a framework to enhance child-robot interaction based on principles of behavior analysis. The system's architecture features an innovative framework capable of online replanning to adapt to user behavior. However, since we have not yet tested the system with children, we predict certain limitations that are likely to emerge and point out directions for future research. ABA-based interventions are often criticized for focusing on behavior modification rather than supporting social acceptance of target behaviors in individuals with neurodevelopmental disorder [3]. For this reason, our model's primary goal is to support behavior management in educational settings, rather than assisting a long-term therapeutic process. Initially, our model will not account for internal and external conditions (e.g. physical discomfort or changes in environment) of the child. Furthermore, we will not evaluate children's acceptance of the robot, their motivation, or attention during the interaction. These factors are crucial for conducting a more precise functional analysis. Future evolution should focus on optimizing our model by incorporating these elements to improve its effectiveness and acceptance.

Acknowledgment

This work has been co-funded by the Italian Ministry of Universities and Research (MUR) under the Piano Nazionale Complementare (PNC), as part of the FIT4MEDROB (PNC0000007) project, Mission 3.

References

[1] J. O. Cooper, T. E. Heron, W. L. Heward, et al., Applied behavior analysis, Pearson/Merrill-Prentice Hall Upper Saddle River, NJ, 2007.

[2] D. A. Wilkenfeld, A. M. McCarthy, Ethical concerns with applied behavior analysis for autism spectrum" disorder", Kennedy Institute of Ethics Journal 30 (2020) 31–69.

[3] J. B. Leaf, J. H. Cihon, R. Leaf, J. McEachin, N. Liu, N. Russell, L. Unumb, S. Shapiro, D. Khosrowshahi, Concerns about aba-based intervention: An evaluation and recommendations, Journal of autism and developmental disorders 52 (2022) 2838–2853.

[4] B. S. Parsonson, Evidence-based classroom behaviour management strategies., Kairaranga 13 (2012) 16–23.

[5] K. B. Geiger, J. E. Carr, L. A. LeBlanc, Function-based treatments for escape-maintained problem behavior: A treatment-selection model for practicing behavior analysts, Behavior Analysis in Practice 3 (2010) 22–32.

[6] F. Alnajjar, C. Bartneck, P. Baxter, T. Belpaeme, M. Cappuccio, C. Di Dio, F. Eyssel, J. Handke, O. Mubin, M. Obaid, et al., Robots in education: An introduction to high-tech social agents, intelligent tutors, and curricular tools, Routledge, 2021.

[7] G. Gordon, C. Breazeal, Bayesian active learning-based robot tutor for children's word-reading skills, in: Proceedings of the AAAI conference on artificial intelligence, volume 29, 2015.

[8] D. Leyzberg, S. Spaulding, B. Scassellati, Personalizing robot tutors to individuals' learning differences, in: Proceedings of the 2014 ACM/IEEE international conference on Human-robot interaction, 2014, pp. 423–430.

[9] S.-S. Yun, H. Kim, J. Choi, S.-K. Park, A robot-assisted behavioral intervention system for children with autism spectrum disorders, Robotics and Autonomous Systems 76 (2016) 58–67.

[10] V. Kostrubiec, J. Kruck, Collaborative research project: Developing and testing a robot-assisted intervention for children with autism, Frontiers in Robotics and AI 7 (2020) 37.

[11] W.-Y. G. Louie, J. Korneder, I. Abbas, C. Pawluk, A study on an applied behavior analysis-based robot-mediated listening comprehension intervention for asd, Paladyn, Journal of Behavioral Robotics 12 (2020) 31–46.

[12] A. Triantafyllidis, A. Alexiadis, D. Elmas, G. Gerovasilis, K. Votis, D. Tzovaras, A social robot-based platform for health behavior change toward prevention of childhood obesity, Universal access in the information society 22 (2023) 1405–1415.

[13] I. Baroni, M. Nalin, M. C. Zelati, E. Oleari, A. Sanna, Designing motivational robot: how robots might motivate children to eat fruits and vegetables, in: The 23rd IEEE International Symposium on Robot and Human Interactive Communication, IEEE, 2014, pp. 796–801.

[14] A. Deshmukh, S. K. Babu, R. Unnikrishnan, S. Ramesh, P. Anitha, R. R. Bhavani, Influencing hand-washing behaviour with a social robot: Hri study with school children in rural india, in: 2019

28th IEEE international conference on robot and human interactive communication (RO-MAN), IEEE, 2019, pp. 1–6.
[15] S. Rossi, F. Ferland, A. Tapus, User profiling and behavioral adaptation for hri: A survey, Pattern Recognition Letters 99 (2017) 3–12.
[16] D. Premack, G. Woodruff, Does the chimpanzee have a theory of mind?, Behavioral and brain sciences 1 (1978) 515–526.
[17] A. F. Winfield, Experiments in artificial theory of mind: From safety to story-telling, Frontiers in Robotics and AI 5 (2018) 75.
[18] N. Rabinowitz, F. Perbet, F. Song, C. Zhang, S. A. Eslami, M. Botvinick, Machine theory of mind, in: International conference on machine learning, PMLR, 2018, pp. 4218–4227.
[19] L. Cominelli, D. Mazzei, D. E. De Rossi, Seai: Social emotional artificial intelligence based on damasio's theory of mind, Frontiers in Robotics and AI 5 (2018) 6.
[20] W. Mou, M. Ruocco, D. Zanatto, A. Cangelosi, When would you trust a robot? a study on trust and theory of mind in human-robot interactions, in: 2020 29th IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), IEEE, 2020, pp. 956–962.
[21] M. Romeo, P. E. McKenna, D. A. Robb, G. Rajendran, B. Nesset, A. Cangelosi, H. Hastie, Exploring theory of mind for human-robot collaboration, in: 2022 31st IEEE International Conference on Robot and Human Interactive Communication (RO-MAN), IEEE, 2022, pp. 461–468.
[22] J. Gray, C. Breazeal, Manipulating mental states through physical action: A self-as-simulator approach to choosing physical actions based on mental state outcomes, International Journal of Social Robotics 6 (2014) 315–327.
[23] E. Cipani, K. M. Schock, et al., Functional behavioral assessment, diagnosis, and treatment: A complete system for education and mental health settings, Springer Publishing Company, 2010.
[24] P. Alberto, A. C. Troutman, Applied behavior analysis for teachers, Pearson Upper Saddle River, NJ, 2013.
[25] J. W. Strachan, D. Albergo, G. Borghini, O. Pansardi, E. Scaliti, S. Gupta, K. Saxena, A. Rufo, S. Panzeri, G. Manzi, et al., Testing theory of mind in large language models and humans, Nature Human Behaviour (2024) 1–11.
[26] S. Dettmer, R. L. Simpson, B. S. Myles, J. B. Ganz, The use of visual supports to facilitate transitions of students with autism, Focus on autism and other developmental disabilities 15 (2000) 163–169.
[27] B. J. Weiner, C. C. Lewis, C. Stanick, B. J. Powell, C. N. Dorsey, A. S. Clary, M. H. Boynton, H. Halko, Psychometric assessment of three newly developed implementation outcome measures, Implementation science 12 (2017) 1–12.